

A regional compound Poisson process for hurricane and tropical storm damage

Simon Mak[†]

Georgia Institute of Technology, Atlanta, USA.

Derek Bingham

Simon Fraser University, Burnaby, Canada.

Yi Lu

Simon Fraser University, Burnaby, Canada.

Summary. In light of intense hurricane activity along the U.S. Atlantic coast, attention has turned to understanding both the economic impact and behaviour of these storms. The compound Poisson-lognormal process has been proposed as a model for aggregate storm damage, but does not shed light on regional analysis since storm path data are not used. In this paper, we propose a fully Bayesian regional prediction model which uses conditional autoregressive (CAR) models to account for both storm paths and spatial patterns for storm damage. When fitted to historical data, the analysis from our model both confirms previous findings and reveals new insights on regional storm tendencies. Posterior predictive samples can also be used for pricing regional insurance premiums, which we illustrate using three different risk measures.

Keywords: Spatial statistics; CAR model; autologistic model; Bayesian; MCMC; compound Poisson process; hurricanes and tropical storms.

1. Introduction

In wake of the devastation caused by intense hurricane and tropical storm (HTS) activity, attention has turned to understanding the tendencies and economic impact from these storms. In particular, insurance companies have expressed interest in developing damage prediction models for these catastrophes (e.g., Changnon, 1996; Musulin, 1997), since such models allow for risk assessment and premium pricing of insurance products. However, to price insurance premiums for each localized, or regional, location (hereafter “regional premiums”), models that predict only damage aggregated over all locations (e.g., Kaas et al., 1988; Pai et al., 2006) are of little use, since they do not account for regional variation. We view these two contrasting goals (i.e., prediction of aggregate vs. localized damage) as macro-level and micro-level, respectively. In this paper, a new micro-level model is proposed that incorporates storm path data and, to the best of our knowledge, is the first to provide regional HTS damage forecasts (where regions here are affected U.S. states).

One of the first HTS damage models proposed was the compound Poisson-lognormal model (Katz, 2002). For this model, the total damage incurred up to time t , $X(t)$, is modelled as:

$$X(t) = \sum_{i=1}^{N(t)} Y^{(i)}, \quad (1)$$

[†]Corresponding author

where $\{N(t), t \geq 0\}$ is a Poisson process for storm counts, and $\{Y^{(i)}\}_{i=1}^{\infty}$ are i.i.d. lognormal-distributed monetary damage incurred for each storm, independent of $\{N(t), t \geq 0\}$. This specification is a common model for aggregate loss in certain lines of insurance (e.g., Zuanetti et al., 2006), as well as banking operational risk models (Luo et al., 2009) and total rainfall models (Thompson, 1984). However, since our goal is to develop a micro-level model that incorporates all historical storm data available by location, we cannot use the framework in (1) directly, since it neglects the underlying path from each storm.

El Niño - Southern Oscillation (ENSO) phases, which categorize sea surface temperatures (SST) in the tropic Pacific Ocean, can also have a notable impact on HTS behaviour (e.g., Pielke and Landsea, 1998; Camargo et al., 2007; Pielke, 2009), and should therefore be included as covariates for prediction. Cabras et al. (2011) explored models with various ENSO and seasonal effects for both storm count and damage, and found that, while ENSO and seasonality have a noticeable impact on counts, only ENSO effects appear to impact damage. Unfortunately, since their model is an extension of (1) with covariates, it also cannot serve as a micro-level model. Their model also examines only storm losses above a certain threshold, whereas our interests lie in modeling for the entire loss distribution.

In this paper, we aim to incorporate all available HTS data - storm paths (defined as the set of regional locations hit by a storm), ENSO phase information and regional damage effects, to jointly predict HTS counts and regional damage. For spatial analysis, the proposed model incorporates conditional autoregressive (CAR) models (Besag, 1974), along with *line-of-sight (LOS) connectivity*, which allows storm path information to be borrowed between locations that do not share a common land boundary (i.e., not physically connected). This new notion of connectivity is motivated by the fact that all historical storm paths impact states that are LOS connected, but some are not physically connected. LOS connectivity is therefore crucial for providing realistic storm path predictions. When fit to historical HTS data (Pielke et al., 2008), the proposed model confirms several previous findings, including the strong ENSO impact on storm damage (e.g., Pielke and Landsea, 1998; Katz, 2002) and on seasonal patterns for storm counts (Lu and Zeng, 2012). More importantly, by integrating storm path data, our model also reveals new insights on regional storm behaviour, including ENSO effects on the spatial patterns for both storm path and damage. Regional damage predictions from the proposed model can be used not only by insurers to provide local risk assessment and to price regional premiums, but also by government policymakers to guide decisions on infrastructure and disaster prevention.

The paper is organized as follows. Section 2 gives a brief exploratory analysis of historical HTS data in Pielke et al. (2008). Section 3 provides an overview of Gaussian and autologistic CAR models. Section 4 presents the three components (storm count, path and damage) of the proposed model, and discusses prior specifications and posterior exploration for each component. Section 5 outlines a simulation study, and Section 6 discusses our findings after fitting the proposed model to historical data. We conclude the paper with some future research directions in Section 7.

2. Exploratory analysis of HTS data

2.1. HTS data and ENSO phases

Data on occurrence times, storm paths and normalized economic losses of HTS along the U.S. Gulf and Atlantic coast from 1900-2005 are reported in Pielke et al. (2008). Although




Figure 1: Bar-charts of monthly HTS frequency and proportions for each ENSO phase.

the set of regional locations provided in this dataset is the set of U.S. states, we refer to these as *locations* from here on to provide a more general exposition of the model. Economic loss in this report was defined as “direct losses associated with a hurricane’s impact as determined in the weeks after the event” (this excludes indirect losses such as demand surges, loss mitigation or other longer macroeconomic effects), and we adopt the same definition here. To provide meaningful comparisons between storms occurring at different times, original reported losses were adjusted to 2005 conditions by a normalization method (called PL05) that accounts for changes in inflation, real-wealth-per-capita and population. We use these normalized losses to fit the proposed model, implying that damage forecasts are for 2005 conditions as well.

ENSO phases classify SST variation in the tropic Pacific into three categories: El Niño (warm anomalies), Neutral (no anomalies) and La Niña (cold anomalies). In particular, the National Oceanic and Atmospheric Administration (NOAA) defines a year to be El Niño, Neutral or La Niña if the three-month temperature average in August, September and October (ASO) is above, within or below 0.5°C normal for that period (e.g., Pielke and Landsea, 1998). Although the official hurricane season runs from Jun. 1st to Nov. 30th, only the three months ASO are used in this definition, since over 95% of high-intensity HTS occur within this timeframe (Landsea, 1993). Unfortunately, NOAA’s categorization is only available from 1950 onwards, and replication of their classification for prior years is difficult due to a changing base-period methodology. In order to incorporate reliable ENSO phase data, we are therefore restricted to analyzing storms occurring after 1950. The final dataset consists of 142 storms from 1950-2005: 33 in El Niño years, 66 in Neutral and 43 in La Niña.

2.2. Exploratory data analysis

Here, we make several observations about HTS counts, paths and damage which motivate the proposed model. The average storm count in El Niño, Neutral and La Niña years is 2.06, 2.44 and 3.31, with standard deviations of 1.61, 1.55 and 1.80, respectively, which suggests that colder SST anomalies tend to increase both storm activity and variation. From Figure 1, which plots bar-charts of HTS frequency and proportion for each month, a majority of La Niña and El Niño storms can be seen to occur in ASO, whereas in Neutral years, storms are more evenly spaced over the hurricane season. This gives visual evidence for differing seasonal patterns in each ENSO phase, which we call *ENSO-seasonal count interactions*. Lu and Zeng (2012) also note a similar interaction effect using a three-state

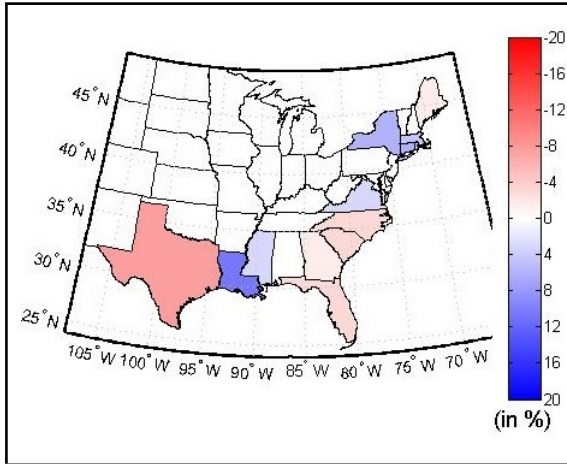


Figure 2a: Difference between El Niño and Neutral hit-rates by location. Blue indicates a higher El Niño rate, and red indicates a higher Neutral one.



Figure 2b: The storm path for Hurricane Gloria (NC, NY, CT, MA) is not spatially connected, but is LOS connected.

hidden Markov model (HMM) for count intensity, with fitted intensities (Figure 6 in their paper) closely resembling continuous analogues of Figure 1. This suggests that the three ENSO phases may indeed be the HMM states, and highlights the need to model ENSO-seasonal interactions.

Next, looking at storm paths in the eastern United States, it turns out that certain locations, e.g., Florida (FL) and Texas (TX), are more often hit than others, e.g., Maine (ME) and Rhode Island (RI), for all three ENSO phases (see Figure A.1 in the Appendix for state map and abbreviations). Indeed, Moran’s I test (Moran, 1950) returns p-values of 0.07, 0.01 and 0.31 against the null hypothesis of no spatial autocorrelation for storm hit-rates in El Niño, Neutral and La Niña, respectively. This provides, at least for the first two phases, some evidence for spatial hit-rate dependencies, and suggests a spatial component is needed for modeling storm paths. Moreover, from Figure 2a, which plots the difference between historical hit-rates in El Niño and Neutral years in each location (with blue indicating higher El Niño hit-rates, and red indicating higher Neutral rates), a subtle ENSO effect is present on spatial patterns as well. Storms in Neutral years concentrate on locations along the southern coast, whereas storms in El Niño tend to move towards northern locations and Louisiana. We call these effects *ENSO-spatial path interactions*.

Several historical storms also hit locations that are not physically connected, i.e., there are states on these paths which do not share a common land boundary with any other locations along the path. For example, the path for Hurricane Gloria, plotted in Figure 2b, hits states that are not physically connected insofar since NC does not share a land boundary with the other three states on the hurricane’s path. Connectedness is an important feature in path prediction, because we do not want to predict storm paths which cannot occur in reality. Consequently, we introduce the key idea of *line-of-sight (LOS) connectedness*: two locations are *LOS neighbours* if there exists an unobstructed sight line from one to the other, and a path is *LOS connected* if it is connected on the graph induced by LOS neighbours. Under this definition, all historical storm paths are LOS connected. This new notion of connectivity is particularly important for the proposed model, since it not only allows for more realistic path predictions, but also permits regional path infor-

Table 1: Median and standard deviation of log-damage for single-location storms at each location and ENSO phase. Number of storms is bracketed.

| | <i>ME</i> | <i>MA</i> | <i>RI</i> | <i>CT</i> | <i>NY</i> | <i>VA</i> | <i>NC</i> |
|--------------------|------------|-------------|-----------|-----------|------------|-----------|--------------|
| El Niño (25) | - | 17.5,NA(1) | - | - | 20.0,NA(1) | - | 15.3,3.0(3) |
| Neutral (58) | 16.5,NA(1) | 19.4,NA(1) | - | - | - | - | 20.0,1.6(8) |
| La Niña (39) | - | 21.8,NA(1) | - | - | - | - | 21.3,1.6(8) |
| Total (122) | 16.5,NA(1) | 19.4,2.2(3) | - | - | 20.0,NA(1) | - | 20.1,1.2(19) |

| | <i>SC</i> | <i>GA</i> | <i>FL</i> | <i>AL</i> | <i>MS</i> | <i>LA</i> | <i>TX</i> |
|--------------------|-------------|-----------|--------------|-----------|-------------|--------------|--------------|
| El Niño (25) | 17.2,2.2(2) | - | 21.2,3.6(8) | - | 18.6,2.2(2) | 19.0,1.4(4) | 16.1,1.8(4) |
| Neutral (58) | 19.0,2.1(5) | - | 19.1,1.8(18) | - | 17.2,NA(1) | 19.2,1.8(8) | 19.3,2.1(16) |
| La Niña (39) | - | - | 20.3,2.6(17) | - | - | 17.5,3.0(4) | 18.8,2.4(9) |
| Total (122) | 18.8,2.3(7) | - | 19.3,2.5(43) | - | 17.2,1.8(3) | 19.1,2.0(16) | 19.0,2.4(29) |

mation to be shared between more locations.

Lastly, for storm damage (measured in log-U.S. \$), La Niña storms incurred the highest median log-damage per storm (20.3), followed by Neutral (19.6) and El Niño (18.8), which is in line with previous findings on the impact of ENSO on storm damage (e.g., Pielke and Landsea, 1998). Spatial patterns for damage are also of interest in our model, since these allow for regional analysis and prediction. Table 1 provides the median and standard deviation of log-damage for the 122 single-location storms, split by location and ENSO (multiple-location storms are not included, since the data only provides aggregated damage over all locations). Although there is some evidence for spatial effects, e.g., FL incurring higher median damage (19.3) than LA (19.1) or TX (19.0), it is difficult to judge if these effects are genuine or due to random error, since many locations contain little to no storm data. We are also interested in *ENSO-spatial damage interactions*, which account for varying spatial patterns of damage for each ENSO phase, but this is even harder to detect from a casual inspection of Table 1, since nearly half of ENSO-location combinations have no data. To estimate these effects, a model incorporating spatial information is needed to borrow data from neighbouring locations.

3. Notation and background

To foreshadow the storm path and damage models, we provide a brief overview of Gaussian and autologistic CAR models; a more detailed discussion can be found in Banerjee et al. (2004). Let $V = \{1, \dots, S\}$ index a set of S areal locations, E be a set of edges indicating spatial connectivity of locations in V , and $G = (V, E)$ be the resulting (undirected) graph. Also, let $\gamma_{\cdot,s} = (\gamma_{1,s}, \gamma_{2,s}, \dots, \gamma_{I,s})^T$ be a vector of I measurements at location s , and let $\gamma = (\gamma_{\cdot,1}^T, \dots, \gamma_{\cdot,S}^T)^T$ be the vector of measurements over all S locations. A *multivariate Gaussian CAR model*, denoted as $MCAR(\Sigma)$, is given by the conditional specification

$$\gamma_{\cdot,s} \mid \gamma_{\cdot,(-s)}, \Sigma \sim N \left(\frac{1}{n_s} \sum_{r \sim s} \gamma_{\cdot,r}, \frac{1}{n_s} \Sigma \right), \quad \forall s = 1, \dots, S, \quad (2)$$

where $\gamma_{\cdot,(-s)} \equiv \{\gamma_{\cdot,r} : r \neq s\}$ is the set of measurements excluding those at location s , Σ is the unscaled conditional covariance matrix of $\gamma_{\cdot,s}$, and n_s is the number of neighbours at location s . The operator \sim indicates adjacency on G , i.e., $r \sim s$ if and only if the edge

(r, s) is in E . By Brook's Lemma (Brook, 1964), the joint density of (2) becomes

$$p(\gamma) \propto \exp \left\{ -\frac{1}{2} \gamma^T [(\mathbf{D} - \mathbf{W}) \otimes \Sigma^{-1}] \gamma \right\}, \quad (3)$$

where \otimes is the Kronecker product operator, \mathbf{W} is the adjacency matrix of G , and $\mathbf{D} = \text{diag}\{n_s\}$. That is, under the conditional specification (2), γ follows a Gaussian distribution with inverse covariance $(\mathbf{D} - \mathbf{W}) \otimes \Sigma^{-1}$.

MCAR models make for convenient spatial priors in Bayesian models for two reasons. First, the specification (2) often allows for direct sampling in each Gibbs sampling step (Chapter 3 in Banerjee et al., 2004). Second, the computation of the inverse covariance $(\mathbf{D} - \mathbf{W}) \otimes \Sigma^{-1}$ requires only the inverse of an $I \times I$ matrix, meaning posterior exploration can be done much quicker with MCAR priors than with other competing areal models that require larger matrix inversions. An MCAR prior also offers a richer correlation structure than is provided by a set of I independent, univariate CAR models, since it captures spatial cross-correlations between pairs of the I measurements through the off-diagonal parameters in Σ . In this sense, MCAR priors offer two advantages. First, as we illustrate for hurricane damage, these cross-correlations have nice interpretations for many natural phenomena, meaning posterior inference on Σ provides valuable insight on the spatial tendencies of the phenomenon itself. Second, when such cross-correlations are indeed present in data, accounting for it allows spatial information to be shared between different measurements, which can then provide lower prediction errors compared to independent CAR priors. This is illustrated in our simulations in Section 5.

However, a problem with (3) is the singularity of $\mathbf{D} - \mathbf{W}$, which implies that the system is over-parametrized. To maintain identifiability, one solution is to impose zero-sum constraints on the vectors $\gamma_{i,\cdot} = (\gamma_{i,1}, \dots, \gamma_{i,S})^T$, $i = 1, \dots, I$ (e.g., Chapter 3 in Banerjee et al., 2004), giving the so-called *intrinsic* MCAR prior. Letting λ_1 and λ_S be the minimum and maximum eigenvalue of $\mathbf{D}^{-1/2} \mathbf{W} \mathbf{D}^{-1/2}$, respectively, another alternative is to add a propriety parameter $\rho \in (1/\lambda_1, 1/\lambda_S)$ to make $\mathbf{D} - \rho \mathbf{W}$ invertible, giving the *proper* MCAR prior. For posterior exploration and prediction, we choose intrinsic priors for two reasons. First, the intrinsic form allows for improved smoothing in locations with few observations (Neelon et al., 2014), which is important for our application since several locations indeed have few observed storms. Second, as noted in Paciorek (2009), the intrinsic MCAR prior gives an intuitive interpretation of each spatial measurement conditionally centered around its neighbours' averages, whereas the proper prior is both difficult to interpret and yields spatial structures that are difficult to justify as a prior (Wall, 2004). We therefore only use the proper MCAR prior for generating data in our simulation study, since it requires an invertible covariance matrix.

This framework can be extended to model binary variables, and forms the basis of the storm path model in Section 4. Let δ_s be an indicator variable for an event at location s (with $\delta_s = 1$ if the event occurs and 0 otherwise), and let $\boldsymbol{\delta} = (\delta_1, \dots, \delta_S)^T$. An *autologistic CAR model*, denoted as $AL(\boldsymbol{\theta}, \phi)$, is given by the conditional specification

$$\text{logit} \{P(\delta_s = 1 \mid \boldsymbol{\delta}_{(-s)})\} = \mathbf{x}_s^T \boldsymbol{\theta} + \phi \sum_{r \sim s} \delta_r, \quad (4)$$

where $\boldsymbol{\delta}_{(-s)} \equiv \{\delta_r : r \neq s\}$, and \mathbf{x}_s is the covariate vector at location s with coefficients $\boldsymbol{\theta}$. In our storm path model, indicators for each location will be used as covariates to

model both spatial main effects and ENSO-spatial interactions. The parameter $\phi > 0$ quantifies the clustering tendency of events; a larger ϕ implies a higher probability for an event occurring at location s given the same event occurs at its neighbours. Using Brook's Lemma again, the joint distribution of (4) becomes

$$p(\boldsymbol{\delta}) = \frac{\exp \left\{ \boldsymbol{\theta}^T \left(\sum_{s=1}^S \delta_s \mathbf{x}_s \right) + \phi \sum_{r \sim s, r < s} \delta_r \delta_s \right\}}{\sum_{\mathbf{c} \in \{0,1\}^S} \exp \left\{ \boldsymbol{\theta}^T \left(\sum_{s=1}^S c_s \mathbf{x}_s \right) + \phi \sum_{r \sim s, r < s} c_r c_s \right\}}, \quad \boldsymbol{\delta} \in \{0,1\}^S. \quad (5)$$

4. New model: Regional HTS damage prediction

The proposed regional prediction model has three parts: storm count, path and damage-per-storm. Only ENSO, seasonal and spatial effects are considered here (since there are limited storm data), but it is straightforward to generalize the framework to include other predictors as well.

4.1. Model specification

The first component of the proposed model looks at storm counts. Similar to Lu and Zeng (2012), the counts process $\{N(t), t \geq 0\}$ is modelled as a non-homogenous Poisson process (PP) with intensity $\lambda(t)$, where t measures time in years. Let $k(t)$ denote the ENSO phase at time $t \geq 0$, with $k(t) = 1$ indicating El Niño, 2 Neutral, and 3 La Niña. In Section 2, we observed that both ENSO effects and ENSO-seasonal interactions are potentially important, therefore we propose the following model for $\lambda(t)$:

$$\log \lambda(t) = \begin{cases} \beta_{k(t)}^0 + \sum_{p=1}^P \left(u_{k(t)}^{(p)} \sin \left[\frac{2\pi p}{7/12} \left(t - \frac{4}{12} \right) \right] + v_{k(t)}^{(p)} \cos \left[\frac{2\pi p}{7/12} \left(t - \frac{4}{12} \right) \right] \right), & \text{if } \frac{4}{12} < t \leq \frac{11}{12}, \\ -\infty, & \text{otherwise.} \end{cases} \quad (6)$$

The above specification admits a different regression model for $\log \lambda(t)$ in each ENSO phase, with β_k^0 as the intercept, and $u_k^{(p)}$ and $v_k^{(p)}$ parametrizing the amplitude of sinusoidal waves with frequencies $\frac{p}{7/12}$ in ENSO phase k , up to a total of P frequencies. Similar formulations can be found in Stolwijk et al., 1999 and Peng et al., 2005. This allows us to model not only the ENSO main effects (i.e., varying count frequencies by ENSO) through different choices of β_k^0 , but also ENSO-seasonal interactions (i.e., varying seasonal patterns by ENSO) through different $u_k^{(p)}$ and $v_k^{(p)}$. Adjusting for these different seasonal patterns is important because it allows the proposed model to account for the lower concentration of storms (and hence storm damage) in ASO for Neutral years (see Figure 2a). Since all storms occurred between the months of May to November inclusive, $\lambda(t)$ is restricted to be positive only within that time period, and zero otherwise. A similar model was proposed in Cabras et al. (2011), however, a key difference is that we model annual seasonality within each ENSO phase.

The second component of our model considers storm paths. Define $G = (V, E)$ as the (undirected) graph with V as the set of locations and E as the set of edges under LOS connectivity. Let $\boldsymbol{\delta} = (\delta_1, \delta_2, \dots, \delta_S)^T$ be the indicator vector for a given storm (with $\delta_s = 1$ if the storm hits location s , and 0 otherwise), and let $\gamma_{k,s}$ be the spatial effect parameter for location s in ENSO phase k . Given a storm count from $\{N(t), t \geq 0\}$ in

ENSO phase k , its path δ is modelled as

$$\text{logit } \{P(\delta_s = 1 \mid \delta_{(-s)})\} = \gamma_k^0 + \gamma_{k,s} + \phi \sum_{r \sim s} \delta_r, \quad \forall s = 1, \dots, S, \quad (7)$$

where \sim is the adjacency operator on G . Each storm path is also assumed to be independent of each other and conditionally independent of $\{N(t), t \geq 0\}$. In light of (4), the specification in (7) allows for different autologistic models in each ENSO phase k , with γ_k^0 as the intercept and $\gamma_{k,\cdot} = (\gamma_{k,1}, \gamma_{k,2}, \dots, \gamma_{k,S})^T$ as the vector of spatial effects. This framework again allows us to model the ENSO main effects (i.e., varying number of locations hit per storm by ENSO) through γ_k^0 , as well as ENSO-spatial interactions (i.e., varying spatial patterns of paths by ENSO) through $\gamma_{k,\cdot}$. To maintain parameter identifiability, we enforce the zero-sum constraints $\sum_{s=1}^S \gamma_{k,s} = 0$, $k = 1, 2, 3$ (Banerjee et al., 2004). The parameter $\phi > 0$ in (7) again quantifies storm clustering tendencies; a larger ϕ increases the chance of a location being hit given its neighbours are also hit.

Although Brook's Lemma can be applied to (7) to get the likelihood (5), we do not recommend doing so, since it assigns positive probability to storm paths that cannot occur (i.e., LOS disconnected paths). Put another way, since all historical storms are LOS connected, allowing for disconnected path predictions results in geographically impossible forecasts. To rectify this, we sum over only LOS connected paths in the denominator of (5), and assign zero probability to LOS disconnected paths (including the null path $\delta = (0, \dots, 0)^T$).

The final component of our model considers regional damage in each location from a given storm. Suppose the i -th storm occurs in ENSO phase k with path $\delta^{(i)} = (\delta_1^{(i)}, \dots, \delta_S^{(i)})^T$, and let $Y_s^{(i)}$ be the damage incurred at location s . If the storm does not hit location s (i.e., $\delta_s^{(i)} = 0$), then clearly $Y_s^{(i)} = 0$. However, if s is indeed hit (i.e., $\delta_s^{(i)} = 1$), then, following Katz (2002), we model $Y_s^{(i)}$ by the lognormal model

$$Y_s^{(i)} \mid \delta^{(i)} \stackrel{\text{indep.}}{\sim} \text{Lognormal} \left(\xi_k^0 + \xi_{k,s} + \zeta^{(i)}, \sigma^2 \right), \quad \forall s : \delta_s^{(i)} = 1, \quad (8)$$

with (8) conditionally independent of both $\{N(t), t \geq 0\}$ and $\delta_{-s}^{(i)}$. The specification in (8) allows for different lognormal models in each ENSO phase, and, as before, ENSO main effects and ENSO-spatial interactions can be modelled by varying ξ_k^0 and $\xi_{k,\cdot} = (\xi_{k,1}, \xi_{k,2}, \dots, \xi_{k,S})^T$. Zero-sum constraints $\sum_{s=1}^S \xi_{k,s} = 0$, $k = 1, 2, 3$ are again enforced for identifiability. The random effect $\zeta^{(i)}$ models storm-to-storm variability for damage (call these *storm severity* effects); larger values of $\zeta^{(i)}$ indicate above-average damage in all affected locations along the path for storm i , and vice versa. We assume that these random effects are i.i.d. $N(0, \sigma_\zeta^2)$.

Lastly, the proposed process for $X_s(t)$, the HTS damage incurred in the time interval $[0, t]$ at location s , is obtained by combining the three submodels:

$$X_s(t) = \sum_{i=1}^{N(t)} \left(\mathbf{1}_{\{\delta_s^{(i)}=1\}} \cdot Y_s^{(i)} \right), \quad s = 1, \dots, S, \quad (9)$$

where $\{N(t), t \geq 0\}$, $\delta^{(i)} \mid \{N(t)\}$ and $Y_s^{(i)} \mid \delta^{(i)}, \{N(t)\}$ are mutually independent. This framework is similar in structure to the compound Poisson-lognormal model (1), with the

Table 2: Hierarchical summary for storm count, path and damage-per-storm submodels.

| <i>Model</i> | <i>Distribution</i> | |
|----------------------------|--|--|
| Count | $[N(t) \lambda(t)]$ | $N(t) \sim \text{Pois}(\lambda(t))$ |
| Intensity | $[\lambda(t) \beta^0, \mathbf{u}, \mathbf{v}]$ | See (6) |
| Priors: | $[\beta^0, \mathbf{u}, \mathbf{v}] = [\beta^0][\mathbf{u}][\mathbf{v}]$ | |
| <i>ENSO</i> | | $[\beta^0] \propto 1$ |
| <i>Seasonality</i> | | $[\mathbf{u}] \propto 1, [\mathbf{v}] \propto 1$ |
| Path | $[\delta^{(i)} \gamma^0, \gamma, \phi, \Sigma_\gamma]$ | See (7) |
| Priors: | $[\gamma^0, \gamma, \phi \Sigma_\gamma] = [\gamma^0][\gamma \Sigma_\gamma][\phi]$ | |
| <i>ENSO</i> | | $[\gamma^0] \propto 1$ |
| <i>Spatial</i> | | $[\gamma \Sigma_\gamma] \sim \text{MCAR}(\Sigma_\gamma)$ |
| <i>Clustering</i> | | $[\phi] \sim \text{IG}(0.01, 0.01)$ |
| Hyperpriors: | $[\Sigma_\gamma]$ | $[\Sigma_\gamma] \sim \text{IW}(4, \mathbf{I})$ |
| Damage | $[Y_s^{(i)} \xi^0, \xi, \zeta, \sigma^2, \sigma_\zeta^2, \Sigma_\xi]$ | See (8) |
| Priors: | $[\xi^0, \xi, \zeta, \sigma^2 \sigma_\zeta^2, \Sigma_\xi] = [\xi^0][\xi \Sigma_\xi][\zeta \sigma_\zeta^2][\sigma^2]$ | |
| <i>ENSO</i> | | $[\xi^0] \propto 1$ |
| <i>Spatial</i> | | $[\xi \Sigma_\xi] \sim \text{MCAR}(\Sigma_\xi)$ |
| <i>Severity</i> | | $[\zeta \sigma_\zeta^2] \sim N(0, \sigma_\zeta^2)$ |
| <i>Log-damage variance</i> | | $[\sigma^2] \sim \text{IG}(0.01, 0.01)$ |
| Hyperpriors: | $[\Sigma_\xi]$ | $[\Sigma_\xi] \sim \text{IW}(4, \mathbf{I})$ |
| | $[\sigma_\zeta^2]$ | $[\sigma_\zeta^2] \sim \text{IG}(0.01, 0.01)$ |

important distinction that (9) allows for regional analysis and damage prediction. Note that, for LOS connected locations r and s , $X_r(t)$ and $X_s(t)$ are correlated through both the storm path and damage submodels. This correlation structure allows for forecasts to be made in locations with little to no observations by borrowing information from LOS neighbours.

4.2. Prior distributions

For statistical inference, we take a Bayesian approach. In some applications, information may exist to guide the selection of prior distributions. Lacking such information, weakly informative or non-informative priors are used instead for parameters in all three submodels. For the spatial effect vectors $\gamma = (\gamma_{\cdot,1}^T, \dots, \gamma_{\cdot,S}^T)^T$ and $\xi = (\xi_{\cdot,1}^T, \dots, \xi_{\cdot,S}^T)^T$, we assign the Gaussian CAR priors $\text{MCAR}(\Sigma_\gamma)$ and $\text{MCAR}(\Sigma_\xi)$, respectively. These priors provide a flexible framework for modeling a wide range of spatial correlation structures. In particular, the diagonal entries in Σ_γ and Σ_ξ account for spatial variation within each ENSO phase, whereas off-diagonal entries control spatial correlation between different pairs of phases. For example, $\Sigma_{\gamma,1,2} > 0$ and $\Sigma_{\xi,1,2} > 0$ indicate that, after accounting for ENSO main effects, the spatial tendencies of storm paths and damage are positively correlated in El Niño and Neutral years. Posterior exploration of Σ_γ and Σ_ξ are therefore not only useful for prediction, but also for gaining a deeper insight on how ENSO influences spatial patterns for both storm path and damage.

As for the remaining model parameters, we assign to each intercept term ($\beta^0 = (\beta_1^0, \beta_2^0, \beta_3^0)^T$ for counts, $\gamma^0 = (\gamma_1^0, \gamma_2^0, \gamma_3^0)^T$ for path, and $\xi^0 = (\xi_1^0, \xi_2^0, \xi_3^0)^T$ for dam-

age) and seasonality parameter ($\mathbf{u} = (\mathbf{u}_{\cdot,1}^T, \dots, \mathbf{u}_{\cdot,P}^T)^T$ and $\mathbf{v} = (\mathbf{v}_{\cdot,1}^T, \dots, \mathbf{v}_{\cdot,P}^T)^T$, where $\mathbf{u}_{\cdot,p} = (u_1^{(p)}, u_2^{(p)}, u_3^{(p)})$ and $\mathbf{v}_{\cdot,p} = (v_1^{(p)}, v_2^{(p)}, v_3^{(p)})$) improper priors on \mathbb{R} . To maintain parameter identifiability, σ^2 , σ_ζ^2 and ϕ are assigned weak inverse-gamma priors $IG(0.01, 0.01)$. The covariance matrices Σ_γ and Σ_ξ are assigned conjugate inverse-Wishart hyperpriors $IW(4, \mathbf{I})$, following Section 3.6 in Gelman et al. (2014). Table 2 provides a concise summary of this prior specification.

4.3. Posterior exploration

For posterior exploration, we use an MCMC hybrid sampler proposed by Tierney (1994), which uses full conditional updates whenever possible (for efficient MCMC mixing) and Metropolis updates (Metropolis et al., 1953) otherwise. Since the three submodels are conditionally independent, separate samplers can be used for each submodel. For the counts submodel, we use the exact likelihood expression for the non-homogenous PP provided in Daley and Vere-Jones (2003), and approximate its integrals by numerical integration. However, under this likelihood, all parameters have non-standard conditional distributions, so the full Metropolis-within-Gibbs sampler (Metropolis et al., 1953) is employed. We briefly describe the samplers for path and damage below, with derivations found in Appendix B.

Let $\Theta^P = \{\gamma^0, \gamma, \phi, \Sigma_\gamma\}$ be the parameter set for the path submodel. After tuning for optimal step-sizes (described later in section) and setting initial parameter values, the path sampler iterates through the following steps:

- Use Metropolis updates for ϕ and each parameter in γ^0 and γ ,
- For $k = 1, 2, 3$, apply zero-sum constraints by centering each $\gamma_{k,\cdot}$ by its mean, as suggested in Banerjee et al. (2004), and
- Update Σ_γ by its full conditional (B.2).

Although sampling $\gamma_{k,\cdot}$ under this “on-the-fly” correction of centering by its mean is not equivalent to sampling directly under the zero-sum constraint (Rue and Held, 2005), the difference is minimal in practice (Paciorek, 2009).

Likewise, let $\Theta^D = \{\xi^0, \xi, \sigma^2, \sigma_\zeta^2, \Sigma_\xi\}$ be the parameter set for the damage submodel. Since $\zeta^{(i)}$ (storm severity effect) and $Y_s^{(i)}$ (regional damage of storm i at location s) are unobserved, both need to be sampled along with the parameters in Θ^D . There are two difficulties in sampling the latter. First, the distribution of independent lognormal random variables conditional on its sum is non-standard (unlike for the gamma distribution, which is proportional to a Dirichlet distribution). Second, the distribution of the sum of lognormals is non-standard as well. To this end, we use an approximation scheme proposed in Mehta et al. (2007) that approximates the sum of lognormals as a new lognormal random variable by matching their respective moment-generating function (mgf) at specific points. Using this approximation, and letting $Y^{(i)}$ be the *observed* total damage from storm i , the conditional regional damage $\{Y_s^{(i)} : \delta_s^{(i)} = 1\} | Y^{(i)}$ can then be sampled with Metropolis updates. Details for this sampling scheme are provided in Appendix C.

After tuning step-sizes and initializing parameter values, the damage sampler cycles through the following steps:

- Use Metropolis steps for each parameter in ξ^0 and ξ ,

- For $k = 1, 2, 3$, apply zero-sum constraints by centering each $\xi_{k,\cdot}$ by its mean,
- Update Σ_ξ and σ_ζ^2 by its full conditionals (B.4) and (B.6),
- Sample $\zeta^{(i)}$ from (B.7) for all storms $i = 1, \dots, n$ in data, and
- For each multiple-location storm, sample $\{Y_s^{(i)} : \delta_s^{(i)} = 1\} | Y^{(i)}$ using the sampling scheme in Appendix C.

For Metropolis steps, a Gaussian random walk proposal distribution is used with step-size equal to its standard deviation. The choice of this step-size is crucial for efficient MCMC mixing, since it allows the sampler to sufficiently explore the parameter space in a reasonable amount of time. An acceptance rate of around 44% is recommended for optimal mixing of univariate MCMC (Roberts et al., 1997), so, to achieve roughly this benchmark for each update in the sampler, we tune step-sizes using the algorithm in Graves (2011). To foreshadow, the acceptance rates in both our simulation and actual analysis ranged from 17% to 49%, and, using the convergence diagnostics described below, the sampled chains show adequate mixing as well.

To assess MCMC convergence, we use trace plot inspection and the Brooks-Gelman-Rubin (BGR) statistic (Gelman et al., 2014), which reports a ratio of within- and between-chain variances (with values much larger than 1 suggesting poor mixing). Close attention is paid to these, since they not only help indicate MCMC convergence, but also help identify parameter identifiability issues (Neelon et al., 2014).

5. Simulation study

We have three goals for our simulation study: to explore MCMC convergence, to ensure good parameter inference, and to assess how well our model provides regional inference and predictions. The last goal is particularly important, since it shows the proposed model can indeed give good regional forecasts with only aggregated damage data. To investigate these goals, several storm datasets were simulated, each incorporating different combinations of ENSO, spatial and seasonal effects for counts, path and damage. Conclusions are consistent for all simulation studies, so, for brevity, we picked only one study to discuss.

The chosen study simulates storm data exhibiting ENSO-seasonal interactions for counts (with $P = 1$), as well as ENSO-spatial interactions for both storm path and damage. The left column in Table 3 summarizes the choice of parameters used for simulation. Each dataset is obtained by first simulating storm counts from (6), then simulating its corresponding path and damage from (7) and (8), respectively. As mentioned in Section 3, to simulate the spatial effects γ and ξ from (3), a parameter $\rho = 0.99$ is added to make $\mathbf{D} - \rho\mathbf{W}$ invertible (this is in contrast to posterior exploration, where the intrinsic form $\mathbf{D} - \mathbf{W}$ is used). To replicate conditions in the original data, the same set of locations and El Niño, Neutral and La Niña years are used. 20,000 MCMC iterations are run for each simulation, with the first 5,000 discarded as burn-in and the remaining samples thinned by keeping every 10-th sample. This procedure is then replicated 100 times to compute average posterior means and coverage rates for highest posterior density (HPD) intervals.

For the first goal, trace plot inspection and BGR statistics show adequate MCMC convergence for a variety of effect combinations. For the second goal, the right part of Table 3 summarizes the average posterior means over all simulations. These posterior means can be seen to be quite close to true values, and coverage rates for 95% HPD

| <i>Parameters</i> | <i>Avg. posterior means</i> |
|---|--|
| $\beta^0 = (1.75, 2, 2.25)$ | $(1.75, 1.97, 2.22)$ |
| $\mathbf{u} = (0, 0.5, -0.5)$ | $(-0.01, 0.51, -0.51)$ |
| $\mathbf{v} = (0, 0.5, -0.5)$ | $(-0.01, 0.52, -0.50)$ |
| $\gamma^0 = (-4.25, -4, -3.75)$ | $(-4.22, -3.90, -3.61)$ |
| $\phi = 1$ | 0.93 |
| $\Sigma_\gamma = 0.2 \cdot \mathbf{I} + 0.8 \cdot \mathbf{J}$ | $\begin{pmatrix} 1.12 & 0.84 & 0.83 \\ 0.84 & 1.15 & 0.85 \\ 0.83 & 0.85 & 1.14 \end{pmatrix}$ |
| $\xi^0 = (18, 20, 22)$ | $(17.8, 20.1, 22.0)$ |
| $\sigma^2 = 5$ | 5.24 |
| $\sigma_\zeta^2 = 1$ | 1.07 |
| $\Sigma_\xi = 0.8 \cdot \mathbf{I} + 0.2 \cdot \mathbf{J}$ | $\begin{pmatrix} 1.32 & 0.08 & 0.10 \\ 0.08 & 1.14 & 0.14 \\ 0.10 & 0.14 & 1.06 \end{pmatrix}$ |

Table 3: Parameter choices and average posterior means for our simulation study. \mathbf{J} is a matrix of ones.




Figure 3: True and pointwise 99% HPD bands (red for MCAR prior, blue for CAR priors) of the lognormal damage density for AL in a Neutral year.

intervals (not shown here) are above 90% for all parameters, indicating that inference for the proposed model (when equipped with weakly-informative priors) performs quite well. This gives confidence that our method can indeed detect effects that are present in the data generating mechanism. Lastly, to assess regional performance, Figure 3 compares the true damage density from (8) and its 99% pointwise HPD bands for a select location (AL) in a randomly chosen simulation, ignoring storm severity. The red band corresponds to the proposed model, which uses an MCAR prior for Σ_ξ , and the blue band corresponds to the model fit using three independent CAR priors. Two interesting observations can be seen. First, the true density is fully contained in the red (MCAR) confidence band, which suggests the proposed model gives good regional inference for damage despite fitting with aggregate data. Second, although both the red and blue bands cover the true density, the latter is noticeably wider than the former. This illustrates an advantage of using MCAR priors: By allowing spatial information to be shared between ENSO phases, the MCAR prior can provide higher precision of estimates compared to independent CAR priors when non-diagonal Σ_γ and Σ_ξ are used to generate the data. This non-diagonality assumption is not unreasonable to make for our application, since, intuitively, one expects locations which are hit often or incur higher damage in one ENSO phase to also be hit often or incur higher damage in other phases as well. For these reasons, we explore only the MCAR model in the following analysis of the HTS data.

6. Analysis of HTS data

6.1. MCMC results and predictions

We now fit the proposed model to the historical HTS data in Section 2. For each submodel, 20,000 iterations of the MCMC sampler are run, with the first 5,000 removed as burn-in. To reduce autocorrelation, the remaining MCMC samples are thinned by keeping every 10-th sample. Figure 4 provides the trace plots (after burn-in and thinning) for a select parameter from each submodel. These plots provide no visual indication of nonstationar-

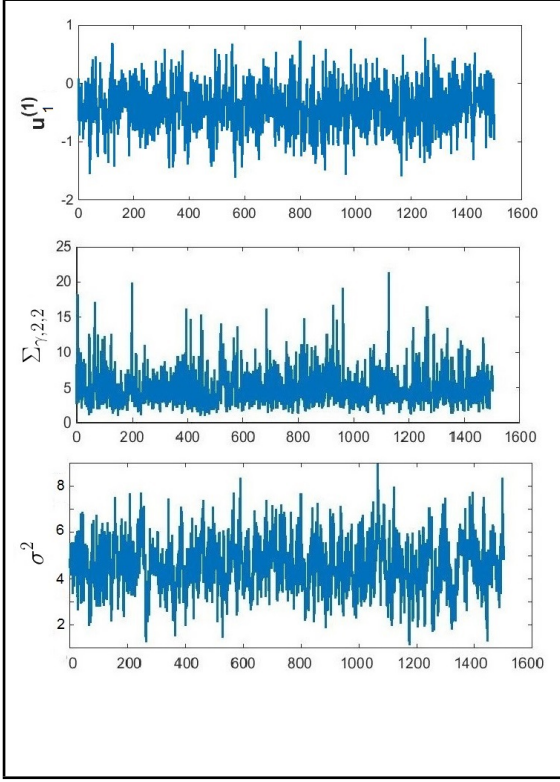


Figure 4: Posterior trace plots of $u_1^{(1)}$ (BGR statistic: 1.0074) from counts model, $\Sigma_{\gamma,2,2}$ (1.0122) from paths model, and σ^2 (0.9982) from damage model.

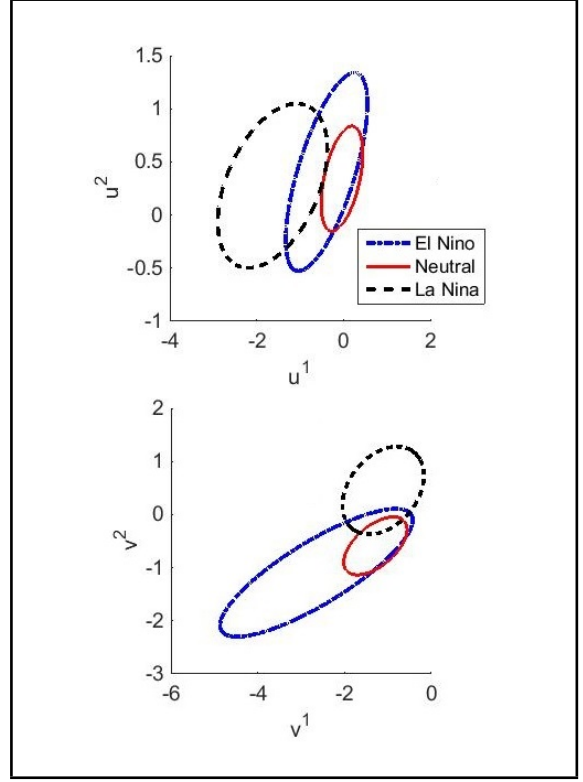


Figure 5: 95% HPD regions for the first two frequency coefficients.

ity and show little autocorrelation as well, which is as desired. This convergence is also confirmed by the BGR statistics, which are all sufficiently close to 1. Similar conclusions hold for remaining parameters as well. Table 4 summarizes the posterior mean and 95% HPD interval for each parameter.

For storm counts, the posterior means and 95% HPD intervals of expected storms in an El Niño, Neutral and La Niña year are 2.06 [1.44, 2.82], 2.44 [1.92, 3.08] and 3.31 [2.37, 4.30], respectively, which confirms previous findings that cooler anomalies are associated with higher storm counts. As for total frequencies P to include, we fit the counts submodel with $P = 1, \dots, 5$ and, using the Deviance Information Criterion (DIC, Spiegelhalter et al., 2002), found that $P = 2$ frequencies provide the best fit for the data. To investigate whether there is evidence for ENSO-seasonal interactions, Figure 5 plots the 95% HPD regions of $(u^{(1)}, u^{(2)})$ and $(v^{(1)}, v^{(2)})$ for each ENSO phase. The regions corresponding to Neutral and La Niña years experience little-to-no overlaps, which suggests that seasonal patterns for these two phases are significantly different, and provides evidence for ENSO-seasonal interactions. On the other hand, the HPD region for Neutral years is nearly contained within that for El Niño, indicating that seasonality for these two phases are quite similar. These results can be seen more clearly in Figure 6, which plots the kernel approximation of historical count intensities with the posterior mean intensities and its 95% pointwise HPD intervals. For El Niño and Neutral years, two intensity peaks can be observed: a large peak at the end of August, and a smaller one at the end of June, which is in line with the two-frequency model selected by DIC. The seasonal pattern in




Figure 6: Estimated historical and posterior mean count intensities for each ENSO phase. The red region marks the 95% pointwise HPD region.

La Niña, however, is visually quite different from the other two ENSO phases, in that it exhibits only the larger peak in August.

As for storm paths, $\hat{\gamma}_1^0 = -4.83$ is larger than $\hat{\gamma}_2^0 = -5.68$ and $\hat{\gamma}_3^0 = -5.91$, where, for the following discussion, $\hat{\gamma}$ denotes the estimated posterior mean of parameter γ . This suggests that more locations tend to be hit in an El Niño storm than in a Neutral or La Niña storm. A strong storm clustering tendency is also implied by the clustering parameter estimate $\hat{\phi} = 1.49$. For spatial variability of storm paths, $\hat{\Sigma}_{\gamma,1,1} = 2.85$, $\hat{\Sigma}_{\gamma,2,2} = 5.97$ and $\hat{\Sigma}_{\gamma,3,3} = 7.30$, which indicates that, although there are noticeable spatial patterns for each ENSO phase, these effects are more pronounced in Neutral and La Niña years than in El Niño. The posterior means for spatial correlations are also quite large, with $\hat{\rho}_{\gamma,1,2} = 0.89$ [0.71, 0.98], $\hat{\rho}_{\gamma,1,3} = 0.89$ [0.70, 0.98] and $\hat{\rho}_{\gamma,2,3} = 0.95$ [0.87, 0.99]. This indicates that locations frequently hit in one ENSO phase tend to be frequently hit in other phases as well.

For storm damage, storms in El Niño years incur notably less damage ($\hat{\xi}_1^0 = 18.60$) than Neutral or La Niña storms ($\hat{\xi}_2^0 = 19.86$, $\hat{\xi}_3^0 = 19.56$), which is consistent with our analysis in Section 2 and findings in Pielke and Landsea (1998). There is also a noticeable storm severity effect, $\hat{\sigma}_{\xi}^2 = 1.66$, although this effect is small compared to the log-damage variance $\hat{\sigma}^2 = 4.75$. For spatial variability of damage, El Niño storms experience the largest variation ($\hat{\Sigma}_{\xi,1,1} = 1.78$), followed by La Niña ($\hat{\Sigma}_{\xi,3,3} = 1.70$) and Neutral storms ($\hat{\Sigma}_{\xi,2,2} = 0.75$). The posterior means and 95% HPD intervals for spatial correlations are $\hat{\rho}_{\xi,1,2} = -0.07$ [-0.36, 0.20], $\hat{\rho}_{\xi,1,3} = -0.17$ [-0.44, 0.07] and $\hat{\rho}_{\xi,2,3} = 0.09$ [-0.21, 0.32], which indicate that the spatial pattern for damage in El Niño storms have a slight negative correlation with that for Neutral and La Niña storms. This is counter-intuitive, since one expects locations that incur higher damage in, say, an El Niño storm, to incur higher damage in, say, a La Niña storm as well. One plausible explanation is that more densely-populated locations are able to better withstand weaker storms, but incur higher damage for more catastrophic storms, which then creates changing spatial patterns for the higher-damaging storms in La Niña and the lower-damaging ones in El Niño.

To further decipher these effects, we generate 1,500 one-year storm forecasts for each ENSO phase from the posterior predictive distribution. For storm paths, Figure 7 plots heat maps of simulated hit-rates in each location. Spatial patterns appear to be quite similar for each ENSO phase, with the two most susceptible locations being FL (simulated

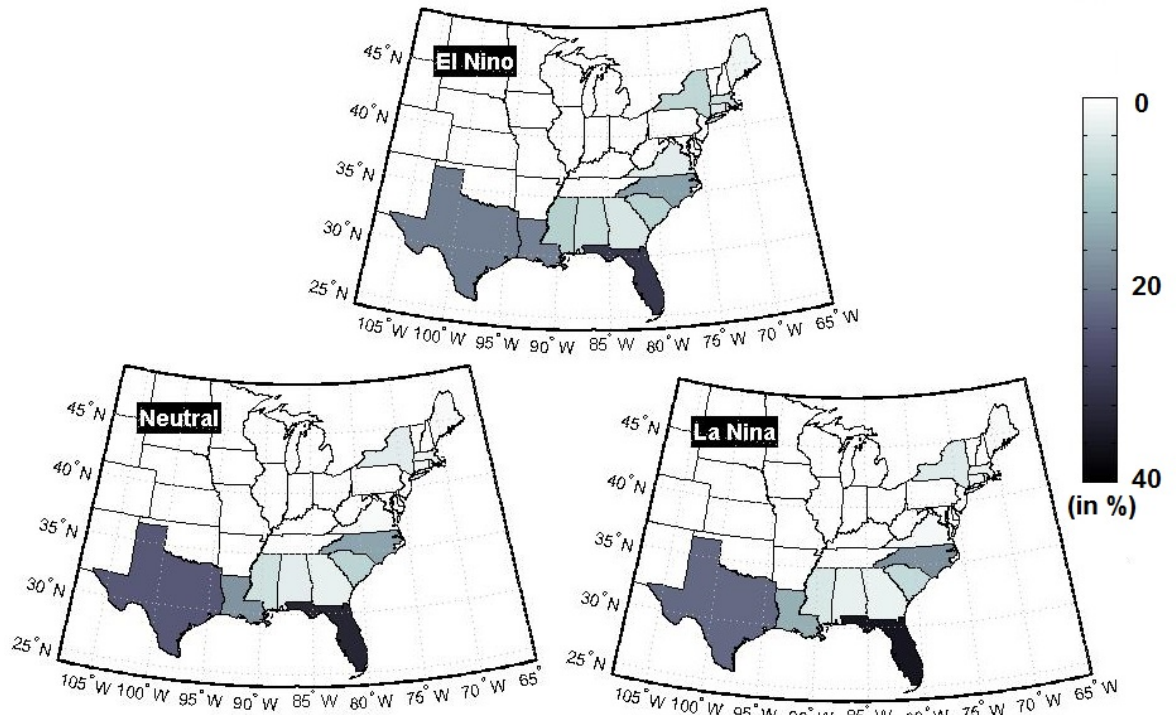


Figure 7: Simulated hit-rates of storms in each ENSO phase.

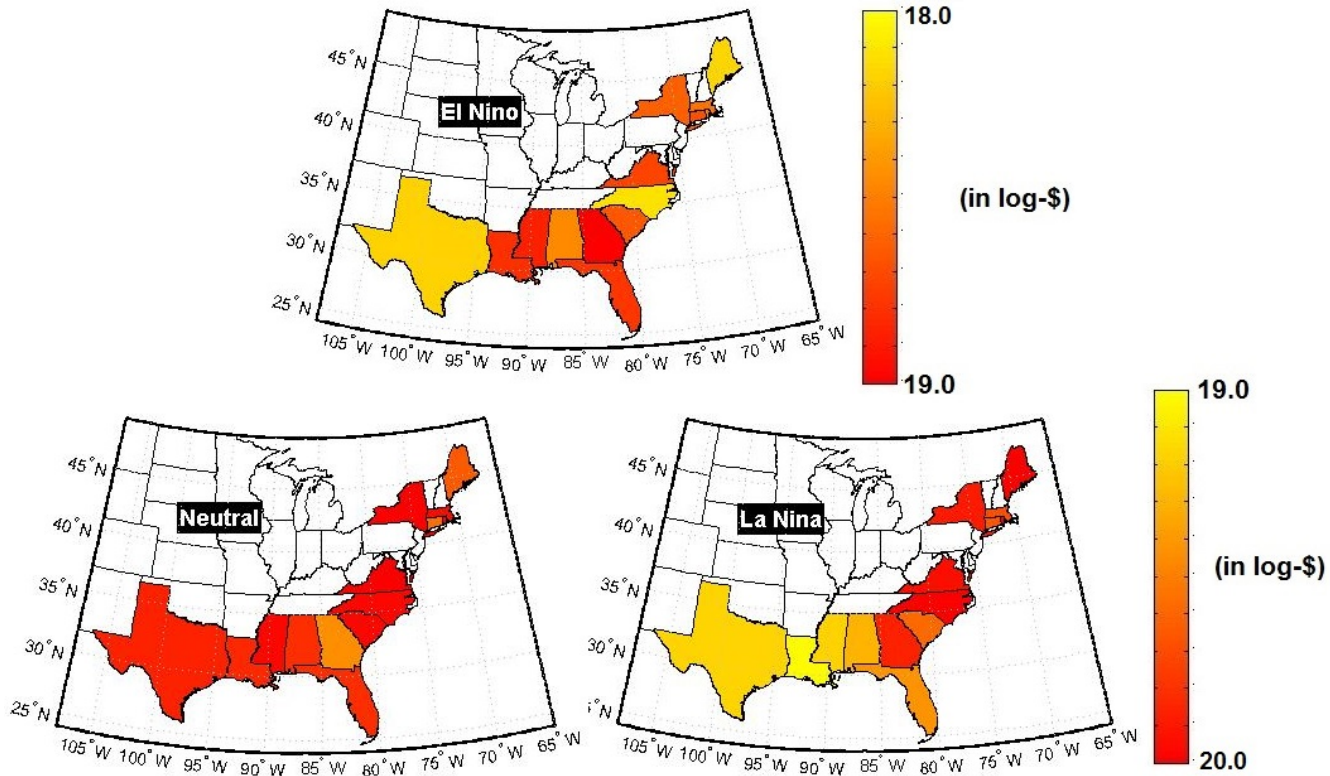


Figure 8: Heat maps for average log-damage of simulated storms in each ENSO phase. A different color scale is used for El Niño to accentuate spatial patterns.

Table 4: Posterior means and 95% HPD intervals for count, path and damage model parameters.

| <i>Model</i> | <i>Parameter</i> | <i>Description</i> | <i>Posterior mean</i> | <i>95% HPD</i> |
|--------------|------------------------|---|-----------------------|--------------------------------|
| Count | β_1^0 | <i>EN</i> main effect | -0.57 | [-1.75, 0.27] |
| | β_2^0 | <i>NE</i> main effect | 0.46 | [0.07, 0.81] |
| | β_3^0 | <i>LN</i> main effect | 0.23 | [-0.52, 0.83] |
| | $u_1^{(1)}, v_1^{(1)}$ | <i>EN</i> seasonality (1 st freq.) | -0.40, -2.60 | [-1.19, 0.36], [-4.5, -1.24] |
| | $u_2^{(1)}, v_2^{(1)}$ | <i>NE</i> seasonality (1 st freq.) | -0.03, -1.29 | [-0.40, 0.36], [-1.92, -0.75] |
| | $u_3^{(1)}, v_3^{(1)}$ | <i>LN</i> seasonality (1 st freq.) | -1.65, -1.12 | [-2.78, -0.72], [-1.97, -0.42] |
| | $u_1^{(2)}, v_1^{(2)}$ | <i>EN</i> seasonality (2 nd freq.) | 0.40, -1.07 | [-0.36, 1.18], [-2.09, -0.25] |
| | $u_2^{(2)}, v_2^{(2)}$ | <i>NE</i> seasonality (2 nd freq.) | 0.35, -0.59 | [-0.04, 0.75], [-1.05, -0.16] |
| | $u_3^{(2)}, v_3^{(2)}$ | <i>LN</i> seasonality (2 nd freq.) | 0.27, 0.45 | [-0.36, 0.88], [-0.19, 1.14] |
| Path | γ_1^0 | <i>EN</i> main effect | -4.83 | [-6.28, -3.34] |
| | γ_2^0 | <i>NE</i> main effect | -5.68 | [-6.99, -4.26] |
| | γ_3^0 | <i>LN</i> main effect | -5.91 | [-7.33, -4.39] |
| | ϕ | Clustering | 1.49 | [0.48, 2.30] |
| | $\Sigma_{\gamma,1,1}$ | Spatial var. - <i>EN</i> | 2.85 | [0.66, 7.57] |
| | $\Sigma_{\gamma,1,2}$ | Spatial cov. - <i>EN/NE</i> | 3.63 | [1.05, 8.93] |
| | $\Sigma_{\gamma,1,3}$ | Spatial cov. - <i>EN/LN</i> | 4.00 | [1.11, 9.91] |
| | $\Sigma_{\gamma,2,2}$ | Spatial var. - <i>NE</i> | 5.97 | [1.12, 14.13] |
| | $\Sigma_{\gamma,2,3}$ | Spatial cov. - <i>NE/LN</i> | 6.13 | [2.30, 14.23] |
| | $\Sigma_{\gamma,3,3}$ | Spatial var. - <i>LN</i> | 7.30 | [2.36, 18.25] |
| Damage | ξ_1^0 | <i>EN</i> main effect | 18.60 | [17.67, 19.50] |
| | ξ_2^0 | <i>NE</i> main effect | 19.86 | [19.16, 20.53] |
| | ξ_3^0 | <i>LN</i> main effect | 19.56 | [18.72, 20.44] |
| | σ^2 | Log-damage var. | 4.75 | [2.22, 7.42] |
| | σ_ζ^2 | Storm severity | 1.66 | [0.02, 4.33] |
| | $\Sigma_{\xi,1,1}$ | Spatial var. - <i>EN</i> | 1.78 | [0.15, 8.68] |
| | $\Sigma_{\xi,1,2}$ | Spatial cov. - <i>EN/NE</i> | -0.14 | [-2.15, 1.46] |
| | $\Sigma_{\xi,1,3}$ | Spatial cov. - <i>EN/LN</i> | -0.47 | [-4.28, 1.70] |
| | $\Sigma_{\xi,2,2}$ | Spatial var. - <i>NE</i> | 0.75 | [0.13, 2.78] |
| | $\Sigma_{\xi,2,3}$ | Spatial cov. - <i>NE/LN</i> | 0.13 | [-1.54, 2.22] |
| | $\Sigma_{\xi,3,3}$ | Spatial var. - <i>LN</i> | 1.70 | [0.16, 7.85] |

hit-rate 32.8%, averaged over all ENSO phases), TX (21.9%) and LA (16.1%), whereas the least susceptible include ME (1.2%), RI (1.8%) and CT (2.3%). This confirms the high spatial path correlations in Table 4 and our observations in Section 2. Despite this similarity, Figure 7 provides some evidence for ENSO-spatial interactions: El Niño and La Niña storm simulations tend to travel further north (hitting the locations MA, RI, NY and CT 12.0% and 9.0% of the time, respectively) compared to Neutral storms (5.8%), which validates our remarks for Figure 2a.

For storm damage, Figure 8 plots heat maps of average log-damage for all simulated storms (multiple-location storms are included here, since the regional damage amounts $\{Y_s^{(i)}\}$ are available from simulations). We see again that El Niño storms incur notably less damage than both Neutral and La Niña, with the latter two incurring similar damage on average. For spatial effects, a global pattern can be detected over all ENSO phases, with FL incurring higher damage in southern locations, and NY and CT incurring higher in the north. This suggests a *shielding effect*: locations which are often the point of land-

| <i>Location</i> | <i>VaR</i> _{95%} | <i>SV</i> _{0.25} | <i>TVaR</i> _{95%} |
|-----------------|---------------------------|---------------------------|----------------------------|
| <i>ME</i> | 0.0 (0.0%) | 0.2 (0.3%) | 0.9 (0.4%) |
| <i>MA</i> | 0.2 (1.6%) | 3.1 (5.7%) | 8.3 (3.7%) |
| <i>RI</i> | 0.0 (0.0%) | 0.9 (1.6%) | 2.9 (1.3%) |
| <i>CT</i> | 0.0 (0.0%) | 3.4 (6.3%) | 8.4 (3.8%) |
| <i>NY</i> | 0.3 (2.3%) | 5.3 (9.7%) | 17.9 (8.0%) |
| <i>VA</i> | 0.0 (0.1%) | 0.8 (1.4%) | 2.9 (1.3%) |
| <i>NC</i> | 1.2 (9.5%) | 6.3 (11.7%) | 23.5 (10.6%) |
| <i>SC</i> | 0.6 (4.3%) | 0.7 (1.2%) | 4.0 (1.8%) |
| <i>GA</i> | 0.1 (0.5%) | 1.7 (3.1%) | 7.4 (3.3%) |
| <i>FL</i> | 4.9 (37.9%) | 8.3 (15.3%) | 43.5 (19.6%) |
| <i>AL</i> | 0.2 (1.3%) | 1.3 (2.4%) | 5.5 (2.5%) |
| <i>MS</i> | 0.6 (4.4%) | 7.4 (13.8%) | 26.6 (12.0%) |
| <i>LA</i> | 3.2 (24.4%) | 13.2 (24.4%) | 60.5 (27.2%) |
| <i>TX</i> | 1.7 (13.5%) | 1.8 (3.3%) | 10.0 (4.5%) |
| Total | 13.0 (100%) | 54.2 (100%) | 222.5 (100%) |

Table 5: Annual regional premiums in bil. U.S. \$ (% of total in brackets) in an El Niño year using $SV_{0.25}$, $VaR_{95\%}$ and $TVaR_{95\%}$.

fall (FL in the south, NY and CT in the north) receive the brunt of storm damage, while remaining locations are shielded from landfall and incur less damage when storm energy gradually dissipates over land (Tuleya, 1994). Figure 8 also provides visual evidence for ENSO-spatial interactions: El Niño storms incur notably higher damage along the southern coast, most notably in GA (average log-damage = 19.0, 95% HPD prediction interval [13.1, 24.7]) and FL (18.8, [13.9, 24.2]), whereas La Niña storms incur higher damage along the northeast coast, most notably in NC (20.0, [14.8, 25.3]), VA (19.9, [14.8, 26.7]) and NY (19.9, [14.9, 25.1]). Since these northern locations generally have higher population densities, this supports our earlier hypothesis that higher-damaging storms (more prevalent in La Niña) impact densely-populated locations more severely. Lastly, damage in Neutral storms experience much less spatial variation than La Niña in Figure 8, which is as observed in Table 4. This can be explained by the shielding effect: since La Niña storms are more likely to make landfall in the higher density locations in the north, one expects these storms to incur higher damage for northern locations (and lower damage for southern locations) compared to Neutral storms.

Although the above observations are intuitive, they are not at all obvious from the exploratory data analysis in Section 2. By incorporating storm path data and using CAR models as spatial smoothers, our model can extract valuable insights on regional storm behaviour.

6.2. Pricing regional premiums

In actuarial literature, a risk measure is a functional \mathcal{H} mapping a loss random variable X to the non-negative real numbers \mathbb{R}_+ , where \mathcal{H} is “assumed in some way to encapsulate the risk associated with the loss” (Hardy, 2006). Three risk measures to illustrate the regional pricing process:

- **Semi-variance** (SV_η): $\mathcal{H}(X) = \mu_X + \eta\sqrt{\sigma_+^2}$, $\eta \geq 0$, where μ_X is the mean of X and

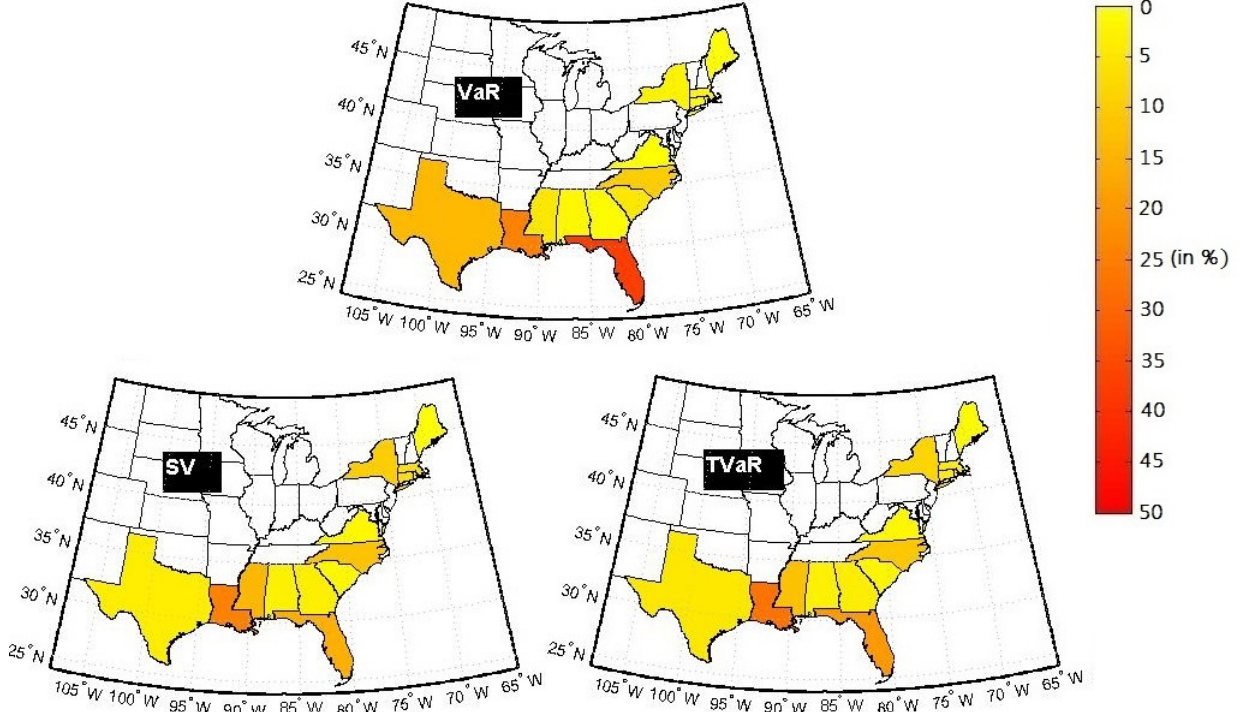


Figure 9: Colour map of premium allocation proportion for $SV_{0.25}$, $VaR_{95\%}$ and $TVaR_{95\%}$.

$\sigma_+^2 = \mathbb{E}[\max(X - \mu_X, 0)^2]$ is the semi-variance of X . SV_η is motivated by the fact that an insurer's risk lies on the right-tail of X , so risk loadings should be calculated on only the positive portion of variance.

- **Value-at-Risk (VaR_α):** $\mathcal{H}(X) = F_X^{-1}(\alpha)$, $0 < \alpha < 1$, where $F_X^{-1}(\alpha)$ is the α -th quantile of X . When an insurer has VaR_α amount of capital available, it incurs a loss with at most probability $1 - \alpha$.
- **Tail-Value-at-Risk ($TVaR_q$):** $\mathcal{H}(X) = \mathbb{E}[X | X > F_X^{-1}(q)]$, $0 < q < 1$. $TVaR_q$ ensures an insurer can cover the average loss given this loss exceeds its q -th percentile.

The parameters η , α and q quantify an insurer's risk-aversion level; larger values imply more risk-averse (and therefore higher) premiums, and vice versa. For illustration, we choose $\eta = 0.25$ and $\alpha = q = 95\%$. We further assume, for simplicity, that there are no deductibles, upper limits or coinsurance, and no discount interest rate for claims as well (these features can easily be added in by modifying MCMC predictions). P&C (Property & Casualty) insurers seldom use VaR pricing in practice, since it, as a risk measure, is incoherent and fails to incorporate important risk information beyond quantiles. However, VaR provides deeper insights on the other two pricing schemes, so we include it in our discussion.

Similar to the previous subsection, the MCMC samples from the full posterior distribution are used to estimate expectations and quantiles for \mathcal{H} . Table 5 summarizes the $SV_{0.25}$, $VaR_{95\%}$ and $TVaR_{95\%}$ premiums for an El Niño year. VaR premiums (totalling \$13.0 bil.) are much lower than the other two measures in all locations, which is expected since VaR pricing ignores heavy right-tailed losses exceeding quantiles. On the other hand, by using conditional excess means to load premiums, $TVaR$ (totalling \$222.5 bil.) over-compensates for heavy-tailed losses, which results in unrealistic pricing schemes.

SV premiums (totalling \$54.2 bil.) appear to strike a good balance between the under-compensation of risk in *VaR* with the over-compensation in *TVaR*.

The spatial allocation of premiums is of primary interest to us, since it highlights the advantages of the proposed model over macro-level models. Figure 9 plots the heat maps for the relative allocation proportion under each risk measure. In all three maps, FL and LA are charged a majority of total premiums, which is expected since these locations have the highest simulated storm hit-rates (Figure 7). *SV* and *TVaR*, however, provide a noticeable spatial smoothing on *VaR*, in that locations with high allocation proportion under *VaR* tend to have slightly lower proportion under *SV* and *TVaR*. There are two reasons for this smoothing. First, although all risk measures incorporate both path effects and damage effects, *VaR* prioritizes path effects by pricing on only quantiles. For example, since FL incurs lower damage-per-storm than NC but has higher hit-rates, it is charged a higher allocation of premiums by *VaR* than by *SV* or *TVaR*. Another reason is that the storm severity effect $\hat{\sigma}_\zeta^2 = 1.66$, which affects all locations along a path, produces similar right-tail loss behaviour for LOS neighbours. For instance, the premium allocation for FL is smoothed much more than for LA, since the former has more LOS neighbours than the latter.

7. Discussion

After a period of intense hurricane activity, most notably Hurricane Katarina in 2005, attention has shifted to understanding the nature and economic impact of these storms. The compound Poisson-lognormal process (1) was proposed for aggregate claims, but cannot serve as a micro-level model since it does not account for underlying storm paths. From an application standpoint, we develop methodology that allows for localized prediction of HTS damage and thus regional pricing of insurance products. A key innovation in the model specification is the use of LOS connectivity rather than physical connectivity, for that not only allows for regional path information to be shared between more locations, but also provides more path predictions that are physically possible. The proposed model could be useful for pricing regional insurance premiums, and possibly for guiding policies on disaster prevention as well.

For posterior exploration of each submodel, a hybrid Metropolis-within-Gibbs sampler is used which exploits direct conditional sampling whenever possible. Simulation studies show adequate MCMC mixing, accurate parameter inference, and good regional predictions using aggregate damage data. When fit to historical data, the proposed model reveals several interesting insights, some new and others confirming previous findings. For storm count, our analysis shows that colder SST anomalies are linked with higher storm counts (as noted in Pielke and Landsea, 1998 and Katz, 2002), and provides evidence for varying seasonal patterns by ENSO (as mentioned in Lu and Zeng, 2012). For storm path, a common spatial pattern is observed over all ENSO phases, with a noticeable higher tendency for El Niño and La Niña storms to hit northern locations. For storm damage, the fitted model supports the association between colder anomalies and higher storm damage, and also reveals a changing spatial pattern for damage, with El Niño storms incurring higher damage in southern locations and La Niña incurring higher damage in the north.

For predictions, we assumed for simplicity that the ENSO phase for the upcoming year

is known. This is not the case in reality: according to IRI ‡, one-year SST forecasts from 25 common forecasts may vary by as much as 1.5 °C! A natural extension of the proposed model is to quantify ENSO phase uncertainty through an appropriate stochastic model. Another future direction is to speed up computation for posterior exploration, since generating 20,000 MCMC samples from historical data takes over 4 hours on a quad-core Ivy Bridge 2.7 GHz processor desktop (with code written in MATLAB version 8.1). After a profile inspection, the bottleneck step is computing the denominator of (5), which requires summing over all LOS connected paths. Several approximations have been proposed (e.g., Hoeting et al., 2000 and Hughes et al., 2011), and it will be interesting to see whether these computational speed-ups require a substantial tradeoff in sampling accuracy.

Our proposed model (9) can be easily extended to model not only economic loss from natural disasters with an underlying physical path, such as tornados, blizzards and wild-fire, but also in certain weather and climate models. One simply has to identify potential covariates (such as ENSO and seasonality) and model spatial interactions for a suitable subset of these covariates (determined through prior knowledge or exploratory analysis). Lognormal losses may not be appropriate for some applications, and other thinner- or thicker-tailed distributions may be more suitable instead, e.g., gamma or generalized extreme value (GEV) distributions. The framework in (9), however, provides a flexible micro-level model which not only provides regional predictions by allowing information to be borrowed between spatial neighbours, but also reveals useful insights on spatial patterns which would otherwise be lost from only a macro-level analysis.

References

- Banerjee, S., Gelfand, A. E. and Carlin, B. P. (2004) *Hierarchical modeling and analysis for spatial data*. London: Chapman & Hall / CRC.
- Besag, J. (1974) Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society: Series B (Methodological)*, **36**, 192–236.
- Brook, D. (1964) On the distinction between the conditional probability and the joint probability approaches in the specification of nearest-neighbour systems. *Biometrika*, 481–483.
- Cabras, S., Castellanos, M. E. and Gamerman, D. (2011) A default Bayesian approach for regression on extremes. *Statistical Modelling*, **11**, 557–580.
- Camargo, S. J., Barnston, A. G., Klotzbach, P. J. and Landsea, C. W. (2007) Seasonal tropical cyclone forecasts. *WMO Bulletin*, **56**, 297–309.
- Changnon, S. A. (1996) *The great flood of 1993: causes, impacts, and responses*. Boulder, Colorado: Westview Press.
- Daley, D. and Vere-Jones, D. (2003) *An Introduction to the Theory of Point Processes, volume I: Elementary Theory and Methods of Probability and its Applications*. Springer.
- Gelman, A., Carlin, J. B., Stern, H. S. and Rubin, D. B. (2014) *Bayesian data analysis*, vol. 2. Taylor & Francis.
- Graves, T. L. (2011) Automatic step size selection in random walk Metropolis algorithms. *arXiv preprint arXiv:1103.5986*.
- Hardy, M. R. (2006) An introduction to risk measures for actuarial applications. *Society of Actuaries exam notes*.
- Hoeting, J. A., Leecaster, M. and Bowden, D. (2000) An improved model for spatially correlated binary responses. *Journal of Agricultural, Biological, and Environmental Statistics*, **5**, 102–114.
- Hughes, J., Haran, M. and Caragea, P. C. (2011) Autologistic models for binary data on a lattice. *Environmetrics*, **22**, 857–871.

‡<http://iri.columbia.edu/our-expertise/climate/forecasts/enso/current/>

- Kaas, R., Van Heerwaarden, A. and Goovaerts, M. J. (1988) Between individual and collective model for the total claims. *ASTIN Bulletin*, **18**, 169–174.
- Katz, R. W. (2002) Stochastic modeling of hurricane damage. *Journal of Applied Meteorology*, **41**, 754–762.
- Landsea, C. W. (1993) A climatology of intense (or major) atlantic hurricanes. *Monthly Weather Review*, **121**, 1703–1713.
- Lu, Y. and Zeng, L. (2012) A nonhomogeneous Poisson hidden Markov model for claim counts. *ASTIN Bulletin*, **42**, 181–202.
- Luo, X., Shevchenko, P. V. and Donnelly, J. B. (2009) Addressing the impact of data truncation and parameter uncertainty on operational risk estimates. *arXiv preprint arXiv:0904.2910*.
- Mehta, N. B., Wu, J., Molisch, A. F. and Zhang, J. (2007) Approximating a sum of random variables with a lognormal. *Wireless Communications, IEEE Transactions on*, **6**, 2690–2699.
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H. and Teller, E. (1953) Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, **21**, 1087–1092.
- Moran, P. A. (1950) Notes on continuous stochastic phenomena. *Biometrika*, 17–23.
- Musulini, R. T. (1997) Issues in the regulatory acceptance of computer modeling for property insurance ratemaking. *Journal of Insurance Regulation*, **15**, 342–359.
- Neelon, B., Gelfand, A. E. and Miranda, M. L. (2014) A multivariate spatial mixture model for areal data: examining regional differences in standardized test scores. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*.
- Paciorek, C. (2009) Technical vignette 5: Understanding intrinsic gaussian markov random field spatial models, including intrinsic conditional autoregressive models.
- Pai, J. S., Shand, K. J. and Wang, X. (2006) Compound poisson model with covariates: a case study on pet insurance. *North American Actuarial Journal*, **10**, 219–234.
- Peng, R. D., Dominici, F., Pastor-Barriuso, R., Zeger, S. L. and Samet, J. M. (2005) Seasonal analyses of air pollution and mortality in 100 U.S. cities. *American Journal of Epidemiology*, **161**, 585–594.
- Pielke, R. A. (2009) United States hurricane landfalls and damages: can one-to five-year predictions beat climatology? *Environmental Hazards*, **8**, 187–200.
- Pielke, R. A., Gratz, J., Landsea, C. W., Collins, D., Saunders, M. A. and Musulin, R. (2008) Normalized hurricane damage in the United States: 1900 - 2005. *Natural Hazards Review*, **9**, 29–42.
- Pielke, R. A. and Landsea, C. W. (1998) Normalized hurricane damages in the United States: 1925-95. *Weather and Forecasting*, **13**, 621–631.
- Roberts, G. O., Gelman, A. and Gilks, W. R. (1997) Weak convergence and optimal scaling of random walk Metropolis algorithms. *The annals of applied probability*, **7**, 110–120.
- Rue, H. and Held, L. (2005) *Gaussian Markov random fields: theory and applications*. CRC Press.
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P. and Van Der Linde, A. (2002) Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **64**, 583–639.
- Stolwijk, A., Straatman, H. and Zielhuis, G. (1999) Studying seasonality by using sine and cosine functions in regression analysis. *Journal of Epidemiology and Community Health*, **53**, 235–238.
- Thompson, C. (1984) Homogeneity analysis of rainfall series: an application of the use of a realistic rainfall model. *Journal of Climatology*, **4**, 609–619.
- Tierney, L. (1994) Markov chains for exploring posterior distributions. *the Annals of Statistics*, 1701–1728.
- Tuleya, R. E. (1994) Tropical storm development and decay: sensitivity to surface boundary conditions. *Monthly Weather Review*, **122**, 291–304.
- Wall, M. M. (2004) A close look at the spatial structure implied by the car and sar models. *Journal of Statistical Planning and Inference*, **121**, 311–324.
- Zuanetti, D., Diniz, C. and Leite, J. (2006) A lognormal model for insurance claims data. *REVSTAT-Statistical Journal*, **4**, 131–142.

Appendices

A. Map and legend

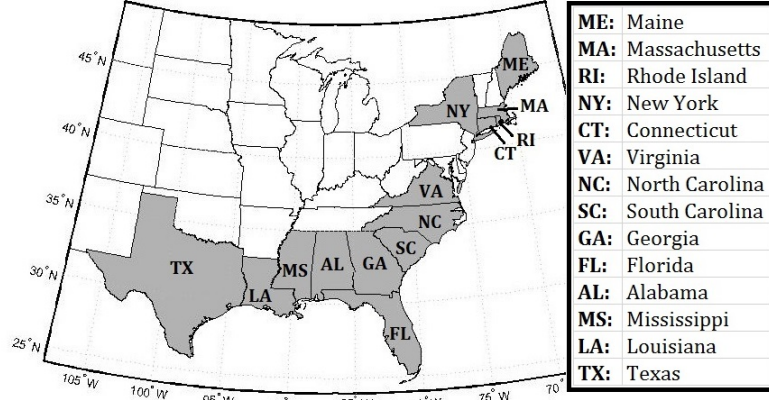


Figure A.1: Map and legend of U.S. states.

B. Derivation of full conditionals

Here, we derive the full conditional distributions of Σ_γ , Σ_ξ , σ^2 and σ_ζ^2 in the storm path and damage samplers. For the path sampler, let $\mathcal{P} = \{\delta^{(i)} : i = 1, \dots, n\}$ be the storm path data for n observed storms. Using the likelihood (3) along with the prior specification in Table 2, the joint posterior of the path submodel parameters Θ^P becomes

$$[\Theta^P | \mathcal{P}] \propto \left[\prod_{i=1}^n p(\delta^{(i)} | \gamma_0, \gamma, \phi) \right] \cdot \left[|\Sigma_\gamma|^{-\text{rank}(\mathbf{D}_w - \mathbf{W})/2} \exp \left\{ -\frac{1}{2} \gamma^T [(\mathbf{D}_w - \mathbf{W}) \otimes \Sigma_\gamma^{-1}] \gamma \right\} \right] \cdot \frac{1}{\phi} \cdot \left[|\Sigma_\gamma|^{-4} \exp \left\{ -\frac{1}{2} \text{tr}(\Sigma_\gamma^{-1}) \right\} \right], \quad (\text{B.1})$$

where $\text{tr}(\cdot)$ is the trace function. Let $\gamma' = (\gamma_{1,\cdot}^T, \gamma_{2,\cdot}^T, \gamma_{3,\cdot}^T)^T$ be a rearranged vector for γ , and, with a slight abuse of notation, let $*$ = $\Theta^P \setminus \theta$, where θ is the current parameter considered. From (B.1), the full conditional distribution of Σ_γ becomes

$$[\Sigma_\gamma | \mathcal{P}, *] \sim IW \left(3 + \text{rank}(\mathbf{D}_w - \mathbf{W}), \gamma'^T (\mathbf{D}_w - \mathbf{W}) \gamma' + \mathbf{I} \right). \quad (\text{B.2})$$

Likewise, for the damage sampler, let $\xi_{\cdot,s} \equiv (\xi_{1,s}, \xi_{2,s}, \xi_{3,s})^T$ for each location $s = 1, \dots, S$, and let $\mathcal{D} = \{\delta^{(i)}, Y_s^{(i)} : i = 1, \dots, n; s = 1, \dots, S\}$ be the set of path data and damage incurred in each location. Using (3) and (8) along with the prior specification in Table 2, the joint posterior for the damage submodel parameters Θ^D and storm severity parameters $\{\zeta^{(i)}\}_{i=1}^n$ becomes

$$[\Theta^D \cup \{\zeta^{(i)}\} | \mathcal{D}] \propto \left[\prod_{i=1}^n \prod_{s: \delta_s^{(i)}=1} \frac{1}{Y_s^{(i)} \sigma} \exp \left\{ -\frac{(\ln Y_s^{(i)} - \xi_{k(t_i)}^0 - \xi_{k(t_i),s} - \zeta^{(i)})^2}{2\sigma^2} \right\} \right] \cdot \left[|\Sigma_\xi|^{-\text{rank}(\mathbf{D}_w - \mathbf{W})/2} \exp \left\{ -\frac{1}{2} \xi^T [(\mathbf{D}_w - \mathbf{W}) \otimes \Sigma_\xi^{-1}] \xi \right\} \right] \cdot \left[\prod_{i=1}^n \frac{1}{\sigma_\zeta} \exp \left\{ -\frac{\zeta^{(i)2}}{2\sigma_\zeta^2} \right\} \right] \cdot \left[(\sigma^2)^{-0.01-1} \exp \left\{ -\frac{0.01}{\sigma^2} \right\} \right] \cdot \left[|\Sigma_\xi|^{-4} \exp \left\{ -\frac{1}{2} \text{tr}(\Sigma_\xi^{-1}) \right\} \right] \cdot (\sigma_\zeta^2)^{-0.01-1} \exp \left\{ -\frac{0.01}{\sigma_\zeta^2} \right\}. \quad (\text{B.3})$$

Let $\boldsymbol{\xi}' = (\boldsymbol{\xi}_{1,\cdot}^T, \boldsymbol{\xi}_{2,\cdot}^T, \boldsymbol{\xi}_{3,\cdot}^T)^T$ be a rearranged vector for $\boldsymbol{\xi}$. From (B.3), the full conditional distribution of $\boldsymbol{\Sigma}_\xi$ becomes

$$[\boldsymbol{\Sigma}_\xi | \mathcal{D}, *] \sim IW \left(4 + \text{rank}(\mathbf{D}_w - \mathbf{W}), \boldsymbol{\xi}'^T (\mathbf{D}_w - \mathbf{W}) \boldsymbol{\xi}' + \mathbf{I} \right), \quad (\text{B.4})$$

the full conditional distribution of σ^2 becomes

$$[\sigma^2 | \mathcal{D}, *] \sim IG \left(\left\{ \sum_{i=1}^n \sum_{s=1}^S \mathbf{1}_{\{\delta_s^{(i)}=1\}} \right\} / 2 + 0.01, \left\{ \sum_{i=1}^n \sum_{s:\delta_s^{(i)}=1} (\ln Y_s^{(i)} - \xi_{k(t_i)} - \xi_{k(t_i),s} - \zeta^{(i)})^2 \right\} / 2 + 0.01 \right). \quad (\text{B.5})$$

the full conditional distribution of σ_ζ^2 becomes

$$[\sigma_\zeta^2 | \mathcal{D}, *] \sim IG \left(\frac{n}{2} + 0.01, \frac{\sum_{i=1}^n \zeta^{(i)2}}{2} + 0.01 \right), \quad (\text{B.6})$$

and the full conditional distributions of storm severities parameters $\{\zeta^{(i)}\}_{i=1}^n$ become

$$[\zeta^{(i)} | \mathcal{D}, *] \sim N \left(\frac{\sigma_\zeta^2}{\sigma_\zeta^2 (\boldsymbol{\delta}^{(i)T} \mathbf{1}) + \sigma^2} \sum_{s:\delta_s^{(i)}=1} (\ln Y_s^{(i)} - \xi_{k(t_i)} - \xi_{k(t_i),s}), \frac{\sigma^2 \sigma_\zeta^2}{\sigma_\zeta^2 (\boldsymbol{\delta}^{(i)T} \mathbf{1}) + \sigma^2} \right). \quad (\text{B.7})$$

C. Conditional sampling of $\{Y_s^{(i)}\}$ given $Y^{(i)}$

Here, we discuss an approximate algorithm for sampling the latent regional damage of a multiple-location storm, given its total incurred damage over all locations. For notational convenience, we consider only one multiple-location storm occurring in ENSO phase k , and drop the superscript i . With a slight abuse of notation, let $1, 2, \dots, M$ be new indices for the set of affected locations $\{s : \delta_s = 1\}$, let Y be the total incurred damage, and let $Y_{m+} = \sum_{l=m+1}^M Y_l$ be the sum of regional damage with new index exceeding m . Following Mehta et al. (2007), we approximate the density of Y_{m+} with a lognormal density having parameters μ_{m+} and σ_{m+}^2 (denoted as $f_{LN}(z; \mu_{m+}, \sigma_{m+}^2)$), with μ_{m+} and σ_{m+}^2 obtained by matching the mgf of Y_{m+} with the mgf of the estimated lognormal distribution at the points -0.001 and -0.005 (suggested in Mehta et al., 2007). The desired distribution can be written as

$$[Y_1, \dots, Y_M | Y, \boldsymbol{\Theta}^D] = [Y_1 | Y, \boldsymbol{\Theta}^D] \cdot [Y_2 | Y_1, Y, \boldsymbol{\Theta}^D] \cdots [Y_{M-1} | Y_1, \dots, Y_{M-2}, Y, \boldsymbol{\Theta}^D] \cdot \mathbf{1} \left\{ \sum_{l=1}^{M-1} Y_l \leq Y \right\}, \quad (\text{C.1})$$

with each term on the right-hand side of (C.1) approximated by:

$$\begin{aligned} [Y_1 = z | Y, \boldsymbol{\Theta}^D] &\propto f_{LN}(z; \xi_k^0 + \xi_{k,1} + \zeta, \sigma^2) \cdot f_{LN}(Y - z; \mu_{1+}, \sigma_{1+}^2) \\ [Y_2 = z | Y_1, Y, \boldsymbol{\Theta}^D] &\propto f_{LN}(z; \xi_k^0 + \xi_{k,2} + \zeta, \sigma^2) \cdot f_{LN}(Y - Y_1 - z; \mu_{2+}, \sigma_{2+}^2) \\ &\vdots \\ [Y_{M-1} = z | Y_1, \dots, Y_{M-2}, Y, \boldsymbol{\Theta}^D] &\propto f_{LN}(z; \xi_k^0 + \xi_{k,M-1} + \zeta, \sigma^2) \cdot \\ &\quad f_{LN} \left(Y - \sum_{l=1}^{M-2} Y_l - z; \mu_{(M-1)+}, \sigma_{(M-1)+}^2 \right). \end{aligned} \quad (\text{C.2})$$

A Metropolis sampler can then be implemented which proposes Y_1, \dots, Y_{M-1} in a single block (rejecting whenever $Y_l < 0$ for some $l = 1, \dots, M-1$ or $\sum_{l=1}^{M-1} Y_l > Y$), and accepts proposals with probabilities computed from (C.1) and (C.2). For better mixing of this sampler, we first scale the sample space of $[Y_1, \dots, Y_M | Y, \Theta^D]$ to the unit simplex, and use an independent $U[0, 1]^{M-1}$ random walk (with step-size tuned using the algorithm in Graves, 2011) for proposals. Note that this algorithm may indeed experience poor mixing when M , the number of locations hit by the same storm, is large. For the HTS data, however, neither M nor the number of spatial locations considered are large ($M \leq 4$ for all storms), so this is not a concern for our application.